



LCG Services Report October – December 2009

6 January 2010

Jamie Shiers

This report covers the final three months of 2009 during which the LHC was successfully restarted and delivered data from proton-proton collisions. It was also the first time that the LHC acted as both an accelerator and a collider: the excitement of this success overrides any small glitches that accompanied it.

Unfortunately service issues remain, ranging from a continued high rate of incidents requiring a “post-mortem” via a Service Incident Report, periodic problems with alarms and other tools / procedures for signaling service problems and a continued high rate of patches / bug fixes for some of the key services. It is possible that this represents the new baseline that we should expect and from which we should perform incremental improvements.

Nonetheless, the WLCG service did deliver with success throughout this quarter, including the Christmas to New Year period, when the experiments ran reprocessing, Monte Carlo production and analysis activities.

This first period of LHC operation also taught us that additional flexibility in rescheduling interventions is required – such as taking advantage of a machine stoppage to carry out an already agreed action.

Globally the service can be considered to be both understood and under control but with further optimizations necessary together with continuity to cover changes already in place – such as in the organization of CERN’s IT department – as well as those foreseen for the first half of 2010, including the end of the EGEE III project and the start of EGI and related activities.

Summary of Main Service Incidents

The following table lists the main service incidents for which a [Service Incident Report](#) was produced. These are typically characterized by a serious degradation or total loss of service of at least several hours. Reporting continues to improve, increasingly using a [template](#). Some sites – including CERN and RAL – also produce reports for their own internal purposes at thresholds lower than those in the WLCG Memorandum of Understanding (MoU). However, this is consistent with the guidelines for producing such reports, as discussed at the WLCG Collaboration workshop held in April 2009 prior to CHEP in Prague:

1. When an MoU target is not met;
2. When requested by the Service Coordinator on Duty (SCOD);
3. When useful for the site / entities own internal purposes.

The reduction in incidents seen in the last quarterly report has not become a trend: there were a total of 15 incidents during Q4 (some lasting days or weeks and arguably multiple incidents) versus 7 in Q3. Thus it is more likely that the effect over the summer was at least in part due to reduced activity (both from the side of the experiments as well as service interventions) as compared to this quarter and prolonged data taking may well

be accompanied by a further increase in such incidents. More work needs to be done on quantifying the severity of incidents – the duration is relatively easy to measure but is clearly not sufficient whereas the recorded impact (e.g. data loss versus performance degradation) is somewhat subjective.

Further details can be found in the weekly reports to the WLCG Management Board and on the WLCG [Service Incident Report](#) page.

Site	When	Issue
PIC	19 Dec	Most of Tier-1 services shutdown to avoid increasing temperature due to cooling failure
IN2P3	08 Dec	Grid services unavailability caused by failure of load balancing mechanism
CERN	02 Dec	Site wide power cut – most CC services down
RAL	30 Nov	LHCb Data Loss Incident at RAL
CERN	20 Nov	SRM / ATLAS high failure rate and restart after thread exhaustion
CERN	18 Nov	CMS Dashboard performance degradation
IN2P3	12 Nov	CMS Data Loss Incident at FR-CCIN2P3
IN2P3	03 Nov	Many services disturbed due to automatic reboot of machines
IN2P3	14 Oct	Batch problem – only very short jobs able to run
CERN	13 Oct	All CASTOR services dead
RAL	09 Oct	Data loss from CASTOR
IN2P3	08 & 10 Oct	SRM service interrupted
RAL	04-09 Oct	CASTOR, LFC and FTS services down
ASGC	27 Sep on	DB problems affecting Grid & later CASTOR services for several weeks

Site Metrics

The metrics listed below continue to provide a simple but rather complete view of whether a site is meeting its service delivery targets.

#	Metric
1	Site is providing (usable) resources that match those pledged & requested;
2	The services are running smoothly, pass the tests and meet reliability and availability targets;
3	“WLCG operations” metrics on handling scheduled and unscheduled service interruptions and degradations are met;
4	Site is meeting or exceeding metrics for “functional blocks”.

Three site reviews were carried out during this quarter, covering ASGC, NL-T1 and RAL. Whilst it took ASGC a large fraction of this quarter to resolve the remaining issues that they had faced over many months, the issues previously raised concerning NL-T1 appear to have been successfully resolved.

The report from the RAL review – organized by GridPP – is still in preparation. In this latter respect it is clear that there have been many improvements since the review held one year earlier. However, there are still staffing concerns in the key areas of data management and databases and the foreseen funding for research in the UK and eventually GridPP4 casts a further shadow.

For all sites follow-up in 2010 will be required, as well as some global actions aimed at optimizing our sharing of knowledge and response to problems as a project / community as a whole. This may be coordinated as part of the EGI InSPIRE “Services for Heavy Users of Distributed Computing Infrastructures” activity and/or the ROSCOE Virtual Research Community (VRC) – aimed at providing end-user and community support, both seeking funding through the EU FP7 programme. Further details are given in the outlook section below.

Outlook for 2010

The daily WLCG operations meetings have proven themselves over a period of 2 years, covering CCRC'08, STEP'09 and the first data taking run of the LHC. Contacts with the LHC operations team have been established and the meeting will have to adapt to changing group structures in CERN's IT – which should be largely transparent to users and external sites. Assuming a positive outcome regarding funding, the role of the EGI InSPIRE work-package, oriented at providing services for “heavy users”, will provide a time limited but much-needed continuation and even increase of service / operations-oriented manpower, whilst the ROSCOE project will have a complementary but more user-oriented focus. The exact work plans for these projects will have to be agreed once the funding levels are known: news is expected on the February / March timeframe.

Working with existing structures, such as those of the WLCG as well as the broader community, e.g. HEPiX, these projects should encourage sharing of tools and techniques to further reduce operations load, increase service reliability and to improve the quality of the service is delivered. This will need to be done not just for WLCG and HEP but also for partner communities involved in these projects, which include astronomy and astrophysics, photon science, life sciences and others. This will not only address the short-medium term issues that can be expected during the “running in” period of the LHC but also – by emphasizing the long-term socio-economic benefits – maximize the chances of future funding for follow-on projects.

Summary and Conclusions

As noted in the report from the previous quarter, the WLCG service continues to deliver at a reasonably reliable and responsive level, with continued improvement as seen on the timescale of months. Well established procedures for responding to exceptions exist and are largely but (still) not always respected. Further improvement will clearly be iterative but is nevertheless required – the “site metric” as described above allows this to be measured quantitatively. A better, but not complete, understanding of how to handle larger service upgrades has been achieved, including the important realization that change is inevitable: it cannot be avoided; it needs to be planned and managed, taking both the long and short term LHC operations schedule into account.